



70-768^{Q&As}

Developing SQL Data Models

Pass Microsoft 70-768 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/70-768.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers



**QUESTION 1**

You are deploying a multidimensional Microsoft SQL Server Analysis Services (SSAS) project. You add two new role-playing dimensions named Picker and Salesperson to the cube. Both of the cube dimensions are based upon the underlying dimension named Employee in the data source view.

Users report that they are unable to differentiate the Salesperson attributes from the Picker attributes.

You need to ensure that the Salesperson and Picker attributes in each dimension use unique names.

In the table below, identify an option that you would use as part of the process to alter the names of the attributes for each of the dimensions.

NOTE: Make only one selection in each column.

Hot Area:

Answer Area

Option	Dimension Picker	Dimension Salesperson
Create a second data source view.	<input type="radio"/>	<input type="radio"/>
Rename the Employee dimension.	<input type="radio"/>	<input type="radio"/>
Create a new named query for both dimensions.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Option	Dimension Picker	Dimension Salesperson
Create a second data source view.	<input type="radio"/>	<input type="radio"/>
Rename the Employee dimension.	<input type="radio"/>	<input type="radio"/>
Create a new named query for both dimensions.	<input checked="" type="radio"/>	<input checked="" type="radio"/>

A named query is a SQL expression represented as a table. In a named query, you can specify an SQL expression to select rows and columns returned from one or more tables in one or more data sources. A named query is like any other

table in a data source view (DSV) with rows and relationships, except that the named query is based on an expression.



A named query lets you extend the relational schema of existing tables in DSV without modifying the underlying data source.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/define-named-queries-in-a-data-source-view-analysis-services>

QUESTION 2

You are developing a SQL Server Analysis Services (SSAS) tabular project. The model includes a table named Sales. The Sales table includes a single date column.

The Sales table must meet the following requirements:

Queries must be able to return all rows.

Must be able to support four different processing schedules for different date ranges.

Date ranges must not include any overlapping data.

You need to implement a solution that meets the requirements.

What should you do?

- A. Create four partitions for the Sales table. Create four roles. Use the same row filter queries for each role and partition.
- B. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for all four tables.
- C. Create four partitions for the Sales table. Use row filter queries for each partition.
- D. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for each of the four tables.

Correct Answer: C

QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have a Microsoft SQL Server Analysis Services (SSAS) multidimensional database that stores customer and order data for customers in the United States only. The database contains the following objects:



Type	Name	Content
Measure	Reseller Average Unit Price	the average unit price of sales
Dimension	Geography	the location of resellers
Hierarchy	Geography.State-Province	the state or province where the reseller is located
Member	Geography.State-Province.&[WA]&[US], Geography.State-Province.&[GA]&[US]	a specific state and country/region



You must create a KPI named Large Sales Target that uses the Traffic Light indicator to display status. The KPI must contain:

Expression type	Description
Value	the reseller average unit price
Goal	the average reseller average unit price for US states other than Colorado (CO)
Status	a green indicator if the value is at least 10 percent above the goal, a red indicator if the value is 15 percent or more below the goal, and a yellow indicator for other values
Trend	the value for trend is always 0



You need to create the KPI.

Solution: You set the value of the Status expression to:

```
AVG({  
  COUSIN(  
    [Geography].[State-Province].&[CO]&[US],  
    [Geography].[State-Province].&[CO]  
  )  
})  
[Measures].[Reseller Average Unit Price])
```



Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 4

You need to configure the CoffeeSale fact table environment. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.



Select and Place:

Actions

Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.

Alter the processing job to run every half during the day.

Alter the client application that queries the cube to query the dimensional data warehouse directly for current day data.

Set the storage mode for all partitions to ROLAP.

Test that the cube meets the functional requirement for data currency and query performance.

Partition the CoffoSale fact table.

Set the storage mode for all partitions to HOLAP.

Alter the processing job to ensure that it rearranges the partition structure each evening.

Answer Area



Correct Answer:



Actions

Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.

Alter the processing job to run every half during the day.

Alter the client application that queries the cube to query the dimensional data warehouse directly for current day data.

Set the storage mode for all partitions to HOLAP.

Answer Area

Partition the CoffeSale fact table.

Set the storage mode for all partitions to ROLAP.

Alter the processing job to ensure that it rearranges the partition structure each evening.

Test that the cube meets the functional requirement for data currency and query performance.

Step 1: Partition the CoffeSale facto table.

Step 2: Set the storage mode for all partitions to HOLAP.

Partitions stored as HOLAP are smaller than the equivalent MOLAP partitions because they do not contain source data and respond faster than ROLAP partitions for queries involving summary data.

Step 3: Alter the processing job to ensure that it rearranges the partition structure each evening.

Step 4: Test that the cube meets the functional requirement for data currency and query performance.

From scenario:

Data analysts must be able to analyze sales for financial years, financial quarters, months, and days. Many reports are based on analyzing sales by month.

The SalesAnalysis cube contains a fact table named CoffeeSale loaded from a table named FactSale in the data warehouse. The time granularity within the cube is 15 minutes. The cube is processed every night at 23:00. You determine that

the fact table cannot be fully processed in the expected time. Users have reported slow query response times.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels-olap-logical-cube-objects/partitions-partition-storage-modes-and-processing>

**QUESTION 5**

You are developing a SQL Server Analysis Services (SSAS) tabular project.

You need to grant the minimum permissions necessary to enable users to query data in a data model.

Which role permission should you use?

- A. Explorer
- B. Process
- C. Browser
- D. Administrator
- E. Select
- F. Read

Correct Answer: F

QUESTION 6

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You have an existing multidimensional cube that provides sales analysis. The users can slice by date, product, location, customer, and employee.

The management team plans to evaluate sales employee performance relative to sales targets. You identify the following metrics for employees:

You need to implement the KPI based on the Status expression.

Solution: You design the following solution:

```
Case
  WHEN ([Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) > 0.9
  THEN 1
  WHEN ([Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) <= 0.9
  AND
  [Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) > 0.74
  THEN 0
  ELSE -1
END
```



Does the solution meet the goal?



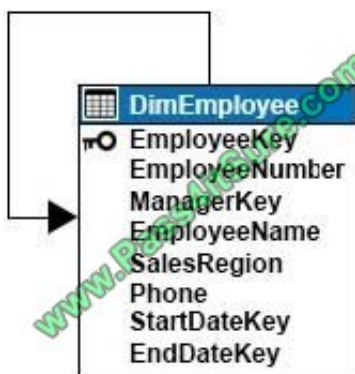
A. Yes

B. No

Correct Answer: B

QUESTION 7

You have a Microsoft SQL Server Analysis Services (SSAS) multidimensional project. You are developing a dimension that uses data from the following table:



The ManagerKey column defines a foreign key constraint that references the EmployeeKey column. The table stores employee history information by using slowly changing dimensions (SCD). Changes to EmployeeName, Phone, or ManagerKey are managed as SCD Type 1 changes. Changes to SalesRegion are managed as SCD Type 2 changes.

You create the following attributes, and set the KeyColumns and NameColumn properties to the columns listed in the table below:

Attribute	KeyColumns	NameColumn
Employee	EmployeeKey	EmployeeName
Employee Number	EmployeeNumber	
Phone	Phone	
Manager	ManagerKey	
Sales Region	SalesRegion	

You need to add a parent-child hierarchy to the dimension to enable navigating the organization hierarchy.

In the table below, identify the attribute that you must use for each attribute usage type.

NOTE: Make only one selection in each column.

Hot Area:



Answer Area

Attribute	Key	Parent
Employee	<input type="radio"/>	<input type="radio"/>
Employee Number	<input type="radio"/>	<input type="radio"/>
Manager	<input type="radio"/>	<input type="radio"/>
Phone	<input type="radio"/>	<input type="radio"/>
Sales Region	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Attribute	Key	Parent
Employee	<input type="radio"/>	<input checked="" type="radio"/>
Employee Number	<input type="radio"/>	<input type="radio"/>
Manager	<input checked="" type="radio"/>	<input type="radio"/>
Phone	<input type="radio"/>	<input type="radio"/>
Sales Region	<input type="radio"/>	<input type="radio"/>

The ManagerKey column, the Manager attribute, defines a foreign key constraint that references the EmployeeKey column, the Employee attribute.

**QUESTION 8**

You need to configure the project option settings to minimize deployment time for the CustomerAnalysis data model. What should you do? To answer, select the appropriate setting from each list in the answer area.

Hot Area:

Answer Area

Location	Setting
Processing option	<div><div></div><div>▼</div></div> <div>Default</div> <div>Do not process</div> <div>Full</div>
Transactional deployment	<div><div></div><div>▼</div></div> <div>False</div> <div>True</div>

Correct Answer:

Answer Area

Location	Setting
Processing option	<div><div></div><div>▼</div></div> <div>Default</div> <div>Do not process</div> <div>Full</div>
Transactional deployment	<div><div></div><div>▼</div></div> <div>False</div> <div>True</div>

Scenario:

Box 1, Processing option:Default

Process Default detects the process state of database objects, and performs processing necessary to deliver unprocessed or partially processed objects to a fully processed state. If you change a data binding, Process Default will do a



Process Full on the affected object.

Note: Processing Method This setting controls whether the deployed objects are processed after deployment and the type of processing that will be performed. There are three processing options:

Default processing (default)

Full processing

None

Box 2, Transactional deployment: False

If this option is False, Analysis Services deploys the metadata changes in a single transaction, and deploys each processing command in its own transaction.

Scenario: The CustomerAnalysis data model will contain a large amount of data and needs to be shared with other developers even if a deployment fails. Each time you deploy a change during development, processing takes a long time.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/deployment-script-files-specifying-processing-options>

QUESTION 9

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series.

Information and details provided in a question apply only to that question.

You are developing Microsoft SQL Server Analysis Services (SSAS) tabular model.

The model must meet the following requirements:

You need to configure model. What should you do?

- A. Ensure that DirectQuery is enabled for the model.
- B. Ensure that DirectQuery is disabled for the model.
- C. Ensure that the Transactional Deployment property is set to True.
- D. Ensure that the Transactional Deployment property is set to False.
- E. Process the model in Process Full mode.
- F. Process the model in Process Data mode.
- G. Process the model in Process Defrag mode.

Correct Answer: A

DAX originally emerged from a Power Pivot add-in for Excel, as a formula language extension for creating calculated columns and measures for data analysis (which are also integral elements of one SSAS Tabular model database, too), but



when Microsoft added support for DAX queries in SQL Server 2012, BI experts started "daxing" data from Tabular model databases.

That trend continues, because of simplicity and fast query execution (related to DirectQuery mode in SSAS Tabular).

References: <https://www.sqlshack.com/query-ssas-tabular-model-database-using-dax-functions/>

QUESTION 10

A database named DB2 uses the InMemory query mode. Users frequently run the following query:

```
EVALUATE
  FILTER (
    ADDCOLUMNS (
      VALUES ('Date' [Calendar Year]),
      "Sales", CALCULATE (SUM ('Internet Sales' [Sales Amount] ) )
    ),
    [Sales] > 8000000
  )
ORDER BY 'Date' [Calendar Year]
```

You need to configure SQL Server Profiler to determine why the query is performing poorly.

Which three event should you monitor on the SQL Server Profiler trace events configuration page? To answer, select the appropriate options in the answer area.

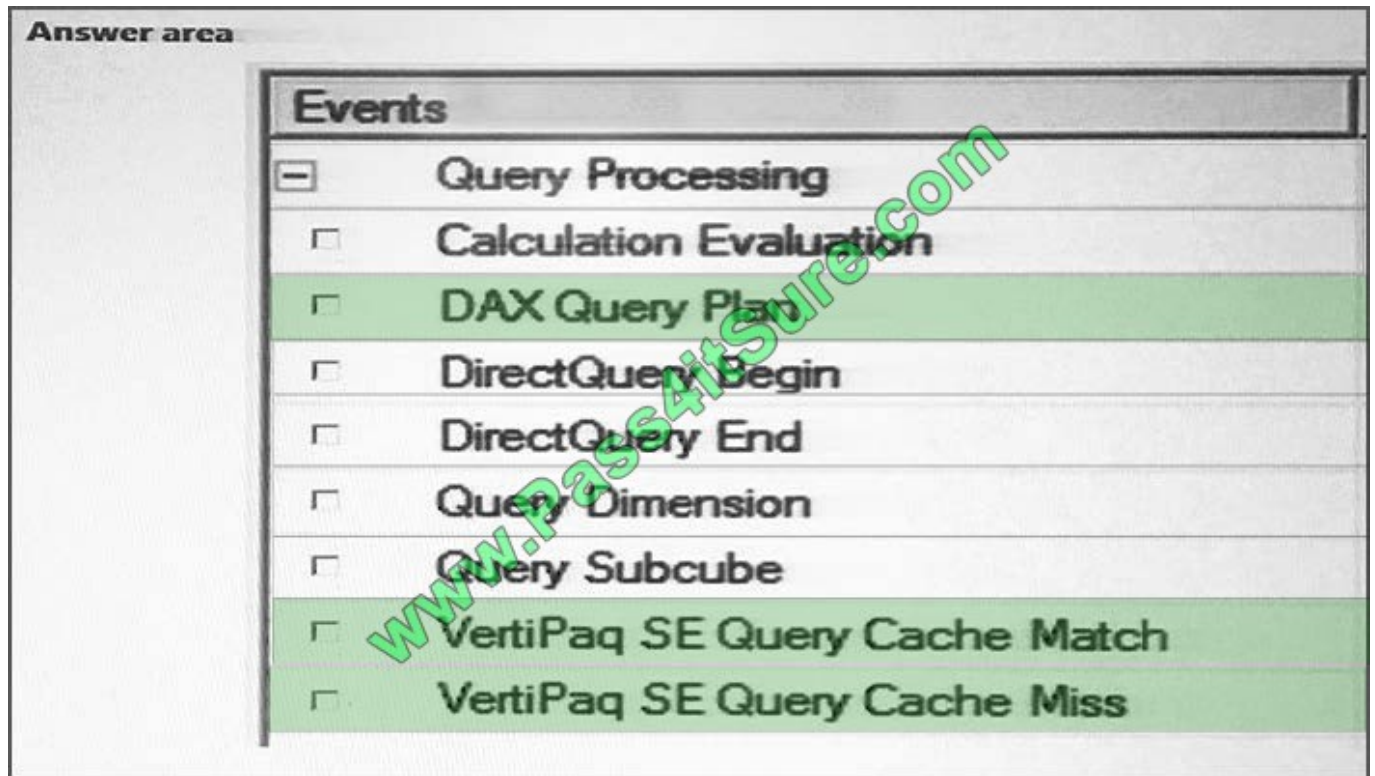
Hot Area:

Answer area

Events	
<input checked="" type="checkbox"/>	Query Processing
<input type="checkbox"/>	Calculation Evaluation
<input type="checkbox"/>	DAX Query Plan
<input type="checkbox"/>	DirectQuery Begin
<input type="checkbox"/>	DirectQuery End
<input type="checkbox"/>	Query Dimension
<input type="checkbox"/>	Query Subcube
<input type="checkbox"/>	VertiPaq SE Query Cache Match
<input type="checkbox"/>	VertiPaq SE Query Cache Miss



Correct Answer:

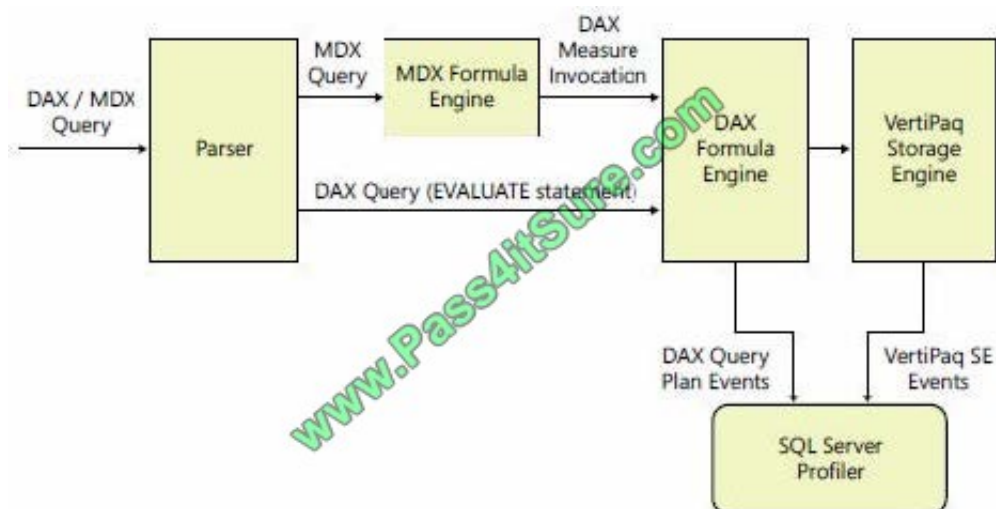


By using SQL Profiler, you can intercept two classes of trace events from Analysis Services, DAX Query Plan and DirectQuery events, both generated by the DirectQuery engine. Here, in this scenario we have a DAX Query.

DAX Query Plan events are generated by the DAX formula.

By using the In-Memory mode, you store a copy of data in the xVelocity (VertiPaq) storage engine.

Figure: This is how a query is executed by using In-Memory mode.



References: Microsoft SQL Server 2012 Analysis Services, The BISM Tabular Model, Microsoft Press (July 2012), page 331 From Scenario: Users report that the query takes a long time to complete.



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/70-768.html>

2022 Latest pass4itsure 70-768 PDF and VCE dumps Download

[70-768 PDF Dumps](#)

[70-768 Study Guide](#)

[70-768 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4itsure.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4itsure, All Rights Reserved.